arXiv:1701.03255v1 [cs.CC] 12 Jan 2017

On the Complexity of L-reachability*

Balagopal Komarath* Jayalal Sarma* K. S. Sunil †

January 13, 2017

Abstract

We initiate a complexity theoretic study of the language based graph reachability problem (L–Reach) : Fix a language L. Given a graph whose edges are labelled with alphabet symbols of the language L and two special vertices s and t, test if there is path P from s to t in the graph such that the concatenation of the symbols seen from s to t in the path P forms a string in the language L. We study variants of this problem with different graph classes and different language classes and obtain complexity theoretic characterizations for all of them. Our main results are the following:

- Restricting the language using formal language theory we show that the complexity of L–Reach increases with the power of the formal language class. We show that there is a regular language for which the L–Reach is NL-complete even for undirected graphs. In the case of linear languages, the complexity of L–Reach does not go beyond the complexity of L itself. Further, there is a deterministic context-free language L for which L–DagReach is LogCFL-complete.
- We use L–Reach as a lens to study structural complexity. In this direction we show that there is a language A in TC^0 for which A–DagReach is NP-complete. Using this we show that P vs NP question is equivalent to P vs DagReach^{-1}(P)[1] question. This leads to the intriguing possibility that by proving DagReach^{-1}(P) is contained in some subclass of P, we can prove an upward translation of separation of complexity classes. Note that we do not know a way to upward translate the separation of complexity classes.

1 Introduction

Reachability problems in mathematical structures are a well-studied problem in space complexity. An important example is the graph reachability problem where given a directed

*A preliminary version of this work with a subset of results was presented at *16th International Workshop on Descriptional Complexity of Formal Systems (DCFS 2014)* and appears in [9].

†Department of Computer Science & Engineering, Indian Institute of Technology Madras, Chennai – 36, India. Email : {baluks,jayalal,sunil}@cse.iitm.ac.in The first author was supported by the TCS Ph.D. Fellowship.

[1]For any complexity class C, DagReach^{-1}(C) = {L : L–DagReach ∈ C}.

graph G and two special vertices s and t, is there a path[2] from s to t in the graph G. This problem exactly captures the space complexity of problems solvable in nondeterministic logarithmic space. Various restrictions of the problem have been studied - reachability in undirected graphs characterizes deterministic logspace [10], reachability in constant width graphs (even undirected) characterizes NC^1 [3], reachability in planar constant width directed graphs characterizes ACC^0 [5] and the version in upward planar constant width directed graphs characterizes AC^0 [4].

A natural extension of the problem using formal language theory is the L–Reach problem: Fix a language L defined over a finite alphabet Σ. Given a graph whose edges are labelled by alphabet symbols and two special vertices s and t, test if there is path from s to t in the graph such that the concatenation of the symbols seen from s to t forms a string in the language L. Indeed, if L is Σ^*, then the string on any path from s to t will be in the language. Hence the problem reduces to the graph reachability problem.

Although L–Reach problem has not been studied from a space complexity theory perspective, a lot is known about its complexity [11, 12, 16, 8, 2]. An immediate observation is that the L–Reach problem is at least as hard as the membership problem of L. Indeed, given a string x, to check for membership in L it suffices to test L–Reach in a simple path of length $|x|$ where the edges are labelled by the symbols in x in that sequence. The literature on the problem is spread over two main themes. One is on restricting the language from the formal language perspective, and the other is by restricting the family of graphs in terms of structure.

An important special case of the problem that was studied is when the language is restricted to be a context-free language (CFL). This is called the CFL–Reach. A primary motivation to study this problem is their application in various practical situations like inter-procedural slicing and inter-procedural data flow analysis [8, 11, 12]. These are used in code optimization, vectorization and parallelization phases of compiler design where one should have information about reaching definitions, available expressions, live variables, etc. associated with the program elements. The goal of inter-procedural analysis is to perform static examination of above properties of a program that consists of multiple procedures. Once a program is represented by its program dependence graph [11], the slicing problem is simply the CFL–Reach problem.

Our Results: The results in this paper are in two flavors.
Results based on Chomsky Hierarchy and Graph Classes: Firstly we study restrictions of L–Reach problem when L is restricted using formal language hierarchy and the graph is restricted to various natural graph classes. Our results on this front are listed in Table 1 (for the sake of completeness, we include some known results too). Apart from the results in Table 1, we show the following theorem for the language class DCFL.

Theorem 1. DCFL–DagReach *is* LogCFL*-complete.*

[2]We follow the convention that a path can have repeated vertices and edges.

Table 1: Formal language class restricted reachability

Language Class	TREE-REACH	DAG-REACH	UREACH/REACH
Regular	Ł-complete[16]	NL-complete[16]	NL-complete (Theorem 4/[16])
Linear	NL-complete (Theorem 5)	NL-complete (Theorem 5)	NL-complete (Theorem 5)
Context-free	LogCFL-complete (Prop. 2)	LogCFL-complete (Prop. 2)	P-complete (Theorem 7/[15])
Context-sensitive	PSPACE-complete (Prop. 3)	PSPACE-complete (Prop. 3)	Undecidable ([2])

Results on the Structural Complexity front: Now we take a complexity theoretic view, where we study L–REACH as an operator on languages. It is shown in Barrett et. al. [2] that even for languages in logspace, the languages L–REACH and L–UREACH are undecidable. Therefore in this section, we consider only DAGs. Note that for any language L, the language L–DAGREACH is decidable.

It is natural to ask whether increasing the complexity of L increases the complexity of L–DAGREACH. More concretely, does $\mathsf{A} \leq_m^{\mathrm{L}} \mathsf{B} \implies \mathsf{A}$–DAGREACH $\leq_m^{\mathrm{L}}$ B–DAGREACH? The following theorem, along with the fact that there exists a language L (see Proposition 2) that is LogCFL-complete for which L–DAGREACH remains LogCFL-complete shows that such a result is highly unlikely.

Theorem 2. *There exists a language* $\mathsf{A} \in \mathsf{TC}^0$ *for which* A–DAGREACH *is* NP*-complete.*

For any complexity class C, we consider the class of languages defined as,

$$\text{DAGREACH}^{-1}(\mathsf{C}) = \{\mathsf{L} : \mathsf{L}\text{–DAGREACH} \in \mathsf{C}\}$$

Note that for any class C, we have $\text{DAGREACH}^{-1}(\mathsf{C}) \subseteq \mathsf{C}$. We have the following theorems for different choices of C.

Theorem : *We show the following structural theorems:*

1. *(Theorem 9)* $\text{DAGREACH}^{-1}(\mathsf{PSPACE}) = \mathsf{PSPACE}$, $\text{DAGREACH}^{-1}(\mathsf{NP}) = \mathsf{NP}$.
2. *(Theorem 10)* $\mathsf{P} \neq \text{DAGREACH}^{-1}(\mathsf{P}) \iff \mathsf{P} \neq \mathsf{NP}$.
3. *(Theorem 11)* $\text{DAGREACH}^{-1}(\mathsf{NL}) \neq \mathsf{NL} \iff \mathsf{NP} \neq \mathsf{NL}$.

The above theorem shows that separating $\text{DAGREACH}^{-1}(\mathsf{P})$ from P would separate P from NP. This gives us an upward translation of lower bounds on complexity classes if we can prove that $\text{DAGREACH}^{-1}(\mathsf{P})$ is contained in some subclass of P. Hence the question whether we can identify some "natural" complexity class containing $\text{DAGREACH}^{-1}(\mathsf{P})$ becomes very interesting. It is clear that $\text{DAGREACH}^{-1}(\mathsf{P})$ contains LogCFL-complete problems but is

highly unlikely to contain some problems in Ł. If $\text{DagReach}^{-1}(\mathsf{P})$ contains some P-complete problem, then proving that $\text{DagReach}^{-1}(\mathsf{P})$ is contained in some subclass of P would be very hard. In this connection, we show the following:

Theorem 3. *If* L *is* P*-complete under syntactic read-once logspace reductions, then* L–DagReach *is* NP*-complete.*

If we are able to extend the above theorem to all types of reductions, then it implies that, assuming NP is not contained in P, $\text{DagReach}^{-1}(\mathsf{P})$ is unlikely to contain P-complete problems. In other words, the above theorem could be interpreted as evidence (albeit very weak evidence) that $\text{DagReach}^{-1}(\mathsf{P})$ may indeed be contained in some subclass of P.

We also remark that Theorem 10 holds with NL instead of P. However, since $\text{DagReach}^{-1}(\mathsf{NL})$ contains NL-complete (under logspace reductions) languages, Theorem 11 is not as promising (as Theorem 10).

A preliminary version of this with a subset of results appears [9]. In [9], we proved that there is a language A in *logspace* such that A–DagReach is NP-complete. In this extended version, we improve this bound to TC^0 (from logspace, see Theorem 2).

2 Preliminaries

In this section, we define language restricted reachability problems and make some observations on their complexity. The definitions for standard complexity classes and their complete problems that we are using in this paper can be found in standard complexity theory textbooks [1]. We use Ł and NL to stand for the complexity classes deterministic logspace and nondeterministic logspace respectively. All reductions (even ones used for defining completeness) in this paper are in logspace unless mentioned otherwise.

Definition 1. *For any language* $\mathsf{L} \subseteq \Sigma^*$*, we consider graph* G *where each edge in* G *is labelled by an element from* Σ*. For any path in* G *we define the yield of the path as the string formed by concatenating the symbols found in the path in that order. Then we define the language* L–Reach *as the set of all* (G, s, t) *such that there exists a path from* s *to* t *in* G *with yield in* L*.*

By restricting the graph in Definition 1, we obtain similar definitions for L–DagReach (DAGs), L–UReach (Undirected Graphs) and L–TreeReach (Orientations of Undirected Trees).

Let Σ and Γ be finite alphabets. A function f from Σ^* to Γ^* is called a projection if for all $x \in \Sigma^*$, the string $f(x) = y$ is such that for all $i \in [m]$, either $y_i = x_j$ for some $j \in [n]$ or $y_i = 0$ or $y_i = 1$, where $m = |y|$ and $n = |x|$. A language L over Σ is said to be projection reducible to a language L' over Γ if there is a projection f such that $x \in L \iff y = f(x) \in \mathsf{L}'$ and $|y|$ is polynomial in $|x|$.

Observation 1. *Any language* L *is projection reducible to* L–TreeReach*.*

Clearly, the above observation holds for any reachability variant based on the graph. This is because L–TreeReach is a restriction of the other reachability variants. In fact the following observation shows that L–TreeReach is not much harder than L.

Observation 2. *For any language* L*, the language* L–TreeReach *is logspace reducible to* L*.*

Observation 2 holds because in logspace we can find the unique path (and hence its yield) from s to t in some tree and run the algorithm for L on the yield.

Next we define classes of languages based on language restricted reachability.

Definition 2. *For any class of languages* C*, we define the set of languages* C–Reach *as the class of all languages* L–Reach *where* L *is in* C*.*

Again, by restricting graphs in Definition 2, we obtain similar definitions for C–DagReach, C–UReach and C–TreeReach.

Definition 3. *For a class of languages* C *and a complexity class* D*, we say that* C–Reach is complete for D *if the following conditions are satisfied.*

- *For all* $\mathsf{L} \in \mathsf{C}$*, the language* L–Reach *is in* D*.*
- *There exists a language* L *in* C *such that the language* L–Reach *is hard for* D*.*

Definition 4. *For any complexity class* C*, we define* $\text{Reach}^{-1}(\mathsf{C})$ *as the set of all languages* L *such that* L–Reach *is in* C*.*

Again, by restricting graphs in Definition 4, we obtain similar definitions for $\text{DagReach}^{-1}(\mathsf{C})$, $\text{UReach}^{-1}(\mathsf{C})$ and $\text{TreeReach}^{-1}(\mathsf{C})$.

Note that by Observation 1, for any class C the relations $\text{Reach}^{-1}(\mathsf{C}) \subseteq \text{DagReach}^{-1}(\mathsf{C}) \subseteq \text{TreeReach}^{-1}(\mathsf{C}) \subseteq \mathsf{C}$ holds. In this paper, we will be mainly studying $\text{DagReach}^{-1}(\mathsf{C})$ for many interesting complexity classes C.

Our motivation in studying $\text{DagReach}^{-1}(\mathsf{C})$ is that it seems that it may be helpful in proving upward translation of separation of complexity classes. Note that we already know, by a standard padding argument, how to translate separations of complexity classes downwards. For example, we know that $\mathsf{NEXP} \neq \mathsf{EXP} \implies \mathsf{P} \neq \mathsf{NP}$. The central question that we address is the following - For a class C, what is the complexity of $\text{DagReach}^{-1}(\mathsf{C})$? Clearly $\text{DagReach}^{-1}(\mathsf{C})$ is contained in C. But for many natural complexity classes C, D and E, if we can show that if $\text{DagReach}^{-1}(\mathsf{C})$ is contained in some subclass D of C, then separating C and D is equivalent to separating C from some complexity class E that contains C.

We use REG, CFL and CSL to stand for well-known formal language classes of regular, context-free and context-sensitive languages respectively[7]. The formal language class LIN, called the set of all linear languages, is the set of all languages with a context-free grammar where the right-hand side of each production consists of at most one non-terminal. The class

LIN can also be characterized as CFLs that can be decided by 1-turn PDAs (sub-family of PDAs where for any computation, the stack height switches only once from non-decreasing mode to non-increasing mode).

We now state a known result with its proof idea which will be used later in the paper.

Proposition 1 ([12])**.** CFL–Reach *is in* P*.*

Proof. (Sketch) The proof is a dynamic programming algorithm. The algorithm maintains for each pair of vertices u and v a table entry $Y[u, v]$ such that $Y[u, v]$ is the set of all non-terminals V in the grammar such that there is a path from u to v with yield that can be derived from V. The algorithm can be modified to output the derivation for x where $x \in \mathsf{L}$ is the yield of a path from s to t. Note that this implies that for all "yes" instances there exists a string with length of the derivation at most polynomial in the size of the graph. □

Sudborough [14] studied the class of languages logspace reducible to a CFL. This class is called LogCFL. Sudborough [14] also showed that LogCFL can be characterized as the set of all languages accepted by an AuxPDA(poly). An AuxPDA(poly) is an NTM with a read-only input tape and a logspace read-write work tape. It also has a pushdown stack available for auxiliary storage. The machine is allowed to run only for a polynomial number (in the input length) of steps. It is also known that the language $\mathsf{NBC}(\mathsf{D}_2)$ (Nondeterministic block choice Dyck_2) is complete for the class LogCFL. The language $\mathsf{NBC}(\mathsf{D}_2)$ consists of all strings of the form $x_1[x_2\#x_3][x_4\#x_5]\dots[x_k\#x_{k+1}]$ where each x_i is a string of two types of parentheses. The string between "[" and "]" is called a block and the symbol # separates choices in a block. A string is in the language $\mathsf{NBC}(\mathsf{D}_2)$ if and only if there is a choice of x_i's from each block such that the final string (after all choices have been made) is in D_2.

3 Formal Language Class restricted Reachability

We know that REG–Reach is in NL [16]. The algorithm works by constructing the product automata of the input graph and the DFA for the regular language. The problem then reduces to the reachability problem on the product automata. One problem with this approach is that even if the input graph is an undirected graph, the product automata will be a directed graph. We know that reachability in directed graphs is harder than reachability in undirected graphs. The following theorem shows that for regular languages, restricted directed and undirected reachability are equivalent.

Theorem 4. *If* L *is the regular language* $\mathsf{L}((\mathsf{ab})^*)$ *over the alphabet* $\{a, b\}$ *then* L–UReach *is* NL*-complete.*

Proof. To show that L–UReach is NL-hard, we give a logspace reduction from REACH. Given an instance (G, s, t) of REACH we construct an instance (G', s, t) of L–UReach where G' is a labelled undirected graph where each edge is labelled either a or b. The vertex set of G' is given by $V(G') = V(G) \cup \{m_{uv} : (u, v) \in E(G)\}$. For each edge $(u, v) \in E(G)$, we add two undirected edges $\{u, m_{uv}\}$ labelled a and $\{m_{uv}, v\}$ labelled b to $E(G')$. It is easy to

see that any directed path from s to t corresponds to a path from s to t in G' labelled by a string in L and vice versa. □

So we know that REG–Reach is NL-complete and CFL–Reach is at least as hard as LogCFL. So it is interesting to consider the complexity of LIN–Reach. We know that $\mathsf{REG} \subseteq \mathsf{LIN} \subseteq \mathsf{CFL}$ in the formal language theory setting. The following theorem shows that LIN–Reach is equivalent to REG–Reach.

Theorem 5. LIN–TreeReach, LIN–DagReach, LIN–UReach *and* LIN–Reach *are all* NL*-complete.*

Proof. There is an NL-complete language in LIN [13]. The hardness follows from this fact and Observation 1. Now we show that all these problems are in NL. The Dynamic Programming algorithm for CFL–Reach from Proposition 1 runs in poly-time and produces a polynomial length derivation for the output string (string yielded by the path). For any language in LIN, a polynomial length derivation can only produce a polynomial length string (and hence polynomial length path). Let us say that the length of the path is bounded by n^k where n is the size of the graph and k is a constant. Then our algorithm will search for a path of length at most n^k by nondeterministically guessing the next vertex at each step and simultaneously parsing the string at each step (using a 1-turn PDA.). This can be implemented by a 1-turn AuxPDA that runs in time n^k and takes $\log(n)$ space. Sudborough [13] proved that this class is exactly the same as NL. □

The following theorem shows that for solving reachability for DCFLs (which are nondeterministic), some nondeterminism is unavoidable.

Theorem 6. DCFL–DagReach *is* LogCFL*-complete.*

Proof. Let $\mathsf{L} \in \mathsf{DCFL}$. We will describe an $\mathsf{AuxPDA}(\mathsf{poly})$ that decides the language L–DagReach. The machine starts with the source vertex s as the current vertex. At each step it nondeterministically moves to an out-neighbor of the current vertex. When the machine takes the edge (u, v) it executes one step of the DPDA for L, using the stack and finite control, with the label on (u, v) as the current input symbol. The machine accepts iff it reaches t and the DPDA accepts.

For hardness, we reduce $\mathsf{NBC}(\mathsf{D}_2)$ to DCFL–DagReach. The reduction results in a series-parallel graph as shown in Figure 1. In the figure, a dashed arrow represents a simple path labelled by the given string. Note that the language D_2 is in DCFL. □

Proposition 2. CFL–TreeReach *and* CFL–DagReach *are* LogCFL*-complete.*

Proof. Sudborough [13] defines a context-free language that is complete for the class LogCFL. This shows the hardness. To show membership in LogCFL consider an $\mathsf{AuxPDA}(\mathsf{poly})$ that starts with s as the current vertex and at each step guesses the next vertex while simultaneously using the stack to simulate the parsing of the CFL. This machine accepts iff the current vertex is t at some point and the PDA is in an accepting state at the same time. It is easy to see that this $\mathsf{AuxPDA}(\mathsf{poly})$ decides these languages. □

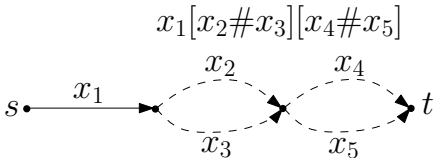

Figure 1: Reducing $\mathsf{NBC}(\mathsf{D}_2)$ to DCFL–DagReach

We now give a simplified presentation of a known result that says that CFL–Reach is P-complete. Also observe that

Theorem 7 ([15]). *Let* D_2 *(ϵ-free Dyck$_2$) be the* CFL *given by the grammar*

$$S \to (S) \mid [S] \mid SS \mid (\ ) \mid [\ ].$$

D_2–Reach *is* P*-complete.*

Proof. This theorem has been proved in [15] using a different terminology. Here we give a simplified presentation of the proof using our terminology for the hardness of this language. We show the P hardness for D_2 by reducing a P-complete problem MCVP (Monotone Circuit Value Problem where fan-out and fan-in of each gate is at most 2) to D_2–Reach. We may assume without loss of generality that each gate in the input circuit has fan-out at most 2. The reduction works by replacing each gate by a gadget as shown in Figure 2. Each gadget in the construction has an input vertex and an output vertex. The gadgets for input gates are straightforward. For an AND gate we add 3 new vertices and connect them to the gadgets for two gates feeding input to the AND gate. Suppose that the left input to the AND gate comes from the 2^{nd} (1^{st}) output wire of the left input gate. Then the first and second edges are labelled by "[" ("(" resp.) and "]" (")" resp.) respectively.

We use proof by induction on the level of the output gate of the circuit to prove the correctness of this reduction. The inductive hypothesis is that there is a valid path from the input vertex to the output vertex of a gadget iff the output of the gate is 1 and any path that enters a gadget through its input gate and leaves it from some vertex other than its output vertex will be invalid. This holds trivially for gadgets for the input gates. Now any valid path from the input vertex to the output vertex of the AND gadget must consist of valid subpaths within the gadgets for the gates feeding input to this AND gate. The only exception is when some path leaves this gadget for the AND gate from some vertex other than its output vertex. Note that by the induction hypothesis such a path can only leave from vertex w or z of the gadget. But the vertex w (also z) has out-degree at most 2 and the other edge will be labelled by a closing bracket that does not match the type of bracket on the edge (u, v). This mismatch invalidates the path. A similar argument holds for OR gates. This completes the induction. □

Now we prove a theorem similar in spirit to Theorem 4 for CFLs. The proof uses the same idea to make the undirected version as hard as the directed one.

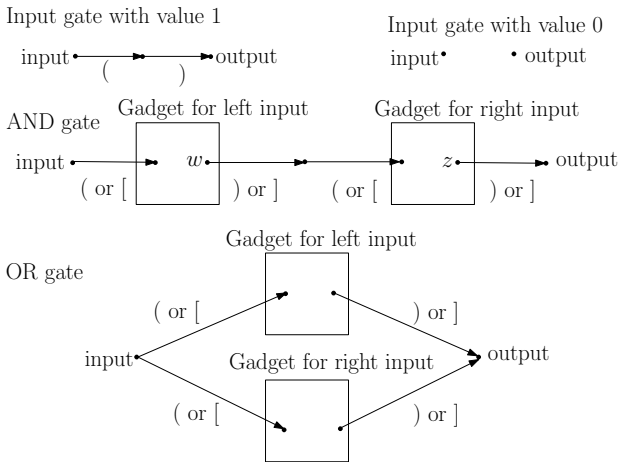

Figure 2: Reducing MCVP to D_2–Reach

Theorem 8. *Let* DD_2 *be the CFL given by the grammar*

$$S \to (aSb) \mid [cSd] \mid SS \mid (a\ b) \mid [c\ d].$$

DD_2–UReach *is* P*-complete.*

Proof. CFL–UReach is in P by [15]. We prove hardness by reducing from D_2–Reach. The reduction works by replacing each edge of the D_2–Reach instance by an undirected path of length two. If for two vertices a, b, the directed edge from a to b is labelled "(" (respectively ")","[" and "]") then replace it by an undirected path of length two with yield "(a"(respectively "b)","[c" and "d]") when read from vertex a to vertex b. The correctness of the reduction is easy to see. □

We state the following proposition, which follows from Theorem 9.

Proposition 3. CSL–TreeReach *and* CSL–DagReach *are* PSPACE*-complete.*

4 Complexity Class restricted Reachability

Now we consider the complexity of L–Reach and its variants when L is chosen from complexity classes. Barrett et. al. [2] has shown that even for languages in Ł, the languages L–Reach and L–UReach are undecidable. But note that for any decidable L, the language L–DagReach is decidable. So we restrict our study only to L–DagReach in this section.

We have seen that moving up in the Chomsky hierarchy increases the complexity of reachability. It is natural to ask whether such an observation also holds with respect to the complexity classes, i.e., increasing the complexity of L increases the complexity of L–DagReach. More concretely, does $\mathsf{A} \leq_m^{\text{Ł}} \mathsf{B}$ imply A–DagReach $\leq_m^{\text{Ł}}$ B–DagReach. The following theorem (which we restate from the introduction) shows that this is very unlikely.

Theorem 2: *There is an* $\mathsf{A} \in \mathsf{TC}^0$ *for which* A–DagReach *is* NP*-complete.*

Proof. The language A can be thought of as an encoding of vertex cover. Each string w in A consists of 3 parts, say w_1, w_2 and w_3. w_1 is a string of the form 1^k0^{n-k} and encodes k, the size of vertex cover, in unary. w_2 consists $\binom{n}{2}$ bits which is the adjacency matrix representation of the input graph. w_3 consists n bits which encodes the vertex cover by the characteristic vector. The strings w_1, w_2 and w_3 are separated by a # and each of the n bits in w_3 is separated by a #.

Let $n_1(x)$ be the number of 1's in the string x. A string w is in the language A iff the following conditions hold.

1. The size of the vertex cover must be at most the size given in the first part of w.
   ie., $n_1(w_3) \leq n_1(w_1)$, and

2. If the edge $\{i, j\}$ is present in the graph, then either the i^{th} or the j^{th} vertex must be present in the vertex cover.
   ie., $(w_2(i, j) = 1) \implies ((w_3(i) = 1) \vee (w_3(j) = 1)$.

Any string $w \in \mathsf{A}$ can be expressed as

$$\begin{aligned}(\forall_{i,j}(w_2(i,j)=1) \quad \implies \quad & (w_3(i)=1 \vee w_3(j)=1)) \wedge \\ & \exists k \leq n, ((n_1(w_1)=k) \wedge (n_1(w_3) \leq k))\end{aligned}$$

An AC^0 circuit is enough to check the conditions $(\forall_{i,j}(w_2(i,j) = 1) \implies (w_3(i) = 1 \vee w_3(j) = 1))$ and $\exists k \leq n, (n_1(w_1) = k)$ but a TC^0 circuit is necessary to check whether $n_1(w_3) \leq k$.

A sketch of the structure of the circuit is given in Fig 3.

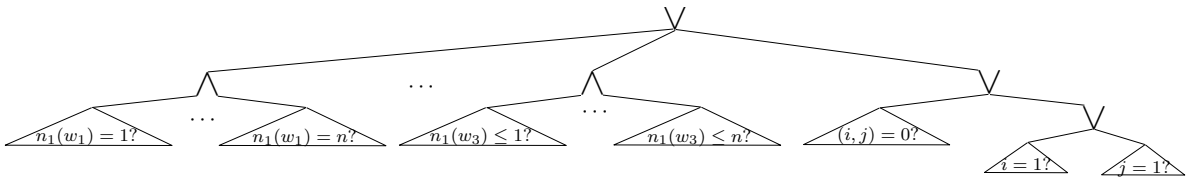

Figure 3: Circuit for A

To show NP-hardness, we reduce $\mathsf{VERTEX\text{–}COVER}$ to A–DagReach.

The language A–DagReach is in NP as the non-deterministic Turing machine guesses the path and verifies whether the yield of the path is in A.

The reduction is given in Fig 4. The DAG contains three parts. The first part, path from s to t_1 encodes the size of the vertex cover and the second part, from t_1 to t_2 encodes the graph while the third part, from t_2 to t represents the actual vertex cover.

For every $w \in \mathsf{A}$ we construct a valid path in DAG as follows. Take the edges labelled by 1, corresponding to the $1's$ in the third part of w (it is same as the vertices in the vertex cover). For the remaining vertices, the edges labelled 0 will be taken in the path.

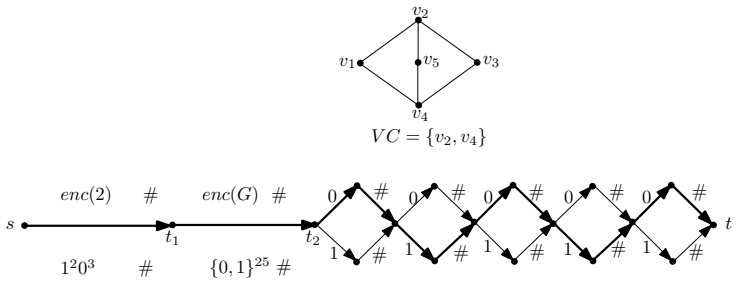

Figure 4: Reducing VERTEX COVER to A–DagReach

Every valid path in the DAG corresponds to a vertex cover in G. Let w be the yield of the path and let w_3 be its third part. Then include i in the vertex cover iff the $(2i-1)^{\text{th}}$ symbol of w_3 is 1. □

We are now going to see how the above result can be used for translating separations of complexity classes upwards (Theorem 10). For any complexity class C, we consider the class of languages defined as $\textsc{DagReach}^{-1}(\mathsf{C}) = \{\mathsf{L} : \mathsf{L}\text{–}\textsc{DagReach} \in \mathsf{C}\}$. We have the following theorems for natural choices of C. Note that for any class C, we have $\textsc{DagReach}^{-1}(\mathsf{C}) \subseteq \mathsf{C}$.

Theorem 9. $\textsc{DagReach}^{-1}(\mathsf{PSPACE}) = \mathsf{PSPACE}$ *and* $\textsc{DagReach}^{-1}(\mathsf{NP}) = \mathsf{NP}$.

Proof. Let $\mathsf{L} \in \mathsf{PSPACE}$, then given an instance of L–DagReach we enumerate all paths from s to t and run the PSPACE algorithm for L on the yield. This is a PSPACE algorithm for L–DagReach. Similarly if $\mathsf{L} \in \mathsf{NP}$, then a path from s to t along with the certificate for the yield on that path is a poly-time verifiable certificate for the L–DagReach problem. □

Theorem 10. $\mathsf{P} \neq \textsc{DagReach}^{-1}(\mathsf{P}) \iff \mathsf{P} \neq \mathsf{NP}$.

Proof. Suppose $\mathsf{P} \neq \textsc{DagReach}^{-1}(\mathsf{P})$ and let $\mathsf{L} \in \mathsf{P} \setminus \textsc{DagReach}^{-1}(\mathsf{P})$. Now L–DagReach is in NP by Theorem 9. By the choice of L we also have L–DagReach is not in P.

For the other direction: suppose $\textsc{DagReach}^{-1}(\mathsf{P}) = \mathsf{P}$. We know that there is a language $\mathsf{L} \in \mathsf{P}$ for which L–DagReach is NP-complete. Hence, $\mathsf{P} = \mathsf{NP}$. □

Theorem 10 shows that separating $\textsc{DagReach}^{-1}(\mathsf{P})$ from P would separate P from NP. This gives us an upward translation of lower bounds on complexity classes provided we can prove that $\textsc{DagReach}^{-1}(\mathsf{P})$ is contained in some subclass of P. The interesting question is whether we can identify some "natural" complexity class containing $\textsc{DagReach}^{-1}(\mathsf{P})$.

By using similar arguments, we also have

Theorem 11. $\textsc{DagReach}^{-1}(\mathsf{NL}) \neq \mathsf{NL} \iff \mathsf{NP} \neq \mathsf{NL}$.

However $\textsc{DagReach}^{-1}(\mathsf{NL})$ contains NL-complete languages (See Theorem 5). So proving that $\textsc{DagReach}^{-1}(\mathsf{NL})$ is separate from NL could be very hard.

The following theorem can be viewed as an evidence that $\textsc{DagReach}^{-1}(\mathsf{P})$ could be separate from P. A language L is syntactic read-once logspace (this notion was considered

by Hartmanis et. al. in [6]) reducible to another language L' iff there is a logspace reduction from L to L' and in the configuration graph for this reduction all paths from the start configuration to the accepting configuration reads each input variable at most once. It shows that if we restrict our attention to syntactic read-once logspace reductions, then L–DagReach for a P-complete problem L is NP-complete. Note that many natural P-complete problems such as CVP (Circuit Value Problem) remains P-complete even under syntactic read-once logspace reductions.

Theorem 3: *If* L *is* P*-complete under syntactic read-once logspace reductions, then* L–DagReach *is* NP*-complete.*

Proof. Let $\mathsf{V} \in \mathsf{NP}$ via a poly-time verifier N. Let W be the witness language for V. i.e., $\mathsf{W} = \{(x, w) : N(x, w) = 1 \text{ and } |w| = |x|^k \text{ for some } k\}$. Since L is P-complete W is read-once logspace reducible to L via M. We reduce V to L–DagReach. Let x be our input. Take the configuration graph G of M on length $|x| + |x|^k$ inputs (after fixing the value of x) and label each edge by the symbol output by the machine M in that step. This graph H is considered as an input to the language L–DagReach. First we prove that $H \in \mathsf{L}$–DagReach implies that $x \in \mathsf{V}$. Consider a path from s to t in H labelled by a string in L. This path corresponds to a witness string for x. Therefore there exists a string w for which (x, w) in W which implies $x \in \mathsf{V}$. For the other direction let $x \in \mathsf{V}$. Therefore there exists a string w such that $(x, w) \in \mathsf{W}$. Now take the path in G that corresponds to this w. The yield of this path is a member of the language L since M outputs this yield when given (x, w) as input. □

5 Discussion and Open Problems

The main result of our work is the observation that if we can prove that the class $\text{DagReach}^{-1}(\mathsf{P})$ is contained in some complexity class that is a subclass of P, then we can translate separation of complexity classes upwards. We propose the following open problem.

Open Problem 1: Prove that $\text{DagReach}^{-1}(\mathsf{P}) \subseteq \mathsf{NC}$.

It would be interesting to study the behavior of $\text{DagReach}^{-1}(.)$ operator on complexity classes below NL. AC^0 is the class of all languages computable by poly-size, constant depth uniform Boolean circuits. Can we say anything about the set of languages $\text{DagReach}^{-1}(\mathsf{AC}^0)$? The only languages L for which we know that L–DagReach is in AC^0 are finite languages. Recall that $\mathsf{DAGREACH}$ is NL-complete and we know that $\mathsf{NL} \neq \mathsf{AC}^0$. Therefore, for any language L such that L–DagReach is in AC^0, the L–DagReach problem is strictly easier than $\mathsf{DAGREACH}$. This leads us to our second open problem.

Open Problem 2: Prove that if L–DagReach $\in \mathsf{AC}^0$ then L is finite.

References

[1] Arora, S., Barak, B.: *Computational Complexity: A Modern Approach*, Cambridge University Press, 2009, ISBN 9780521424264.

[2] Barrett, C. L., Jacob, R., Marathe, M. V.: Formal-Language-Constrained Path Problems, *SIAM Journal of Computing*, **30**(3), 2000, 809–837.

[3] Barrington, D. A. M.: Bounded-Width Polynomial-Size Branching Programs Recognize Exactly Those Languages in NC^1, *Journal of Computer and System Sciences*, **38**(1), 1989, 150–164.

[4] Barrington, D. A. M., Lu, C.-J., Miltersen, P. B., Skyum, S.: Searching constant width mazes captures the AC^0 hierarchy, *In Proceedings of the 15th Annual Symposium on Theoretical Aspects of Computer Science*, Springer-Verlag, 1998.

[5] Hansen, K. A.: Constant Width Planar Computation Characterizes ACC^0, *Proceedings of the 21st Annual Symposium on Theoretical Aspects of Computer Science*, 2004.

[6] Hartmanis, J., Immerman, N., Mahaney, S. R.: One-Way Log-Tape Reductions, *Proceedings of 19th Annual Symposium on Foundations of Computer Science*, 1978.

[7] Hopcroft, J. E., Motwani, R., Ullman, J. D.: *Introduction to automata theory, languages, and computation - international edition (2. ed)*, Addison-Wesley, 2003, ISBN 978-0-321-21029-6.

[8] Horwitz, S., Reps, T. W., Binkley, D.: Interprocedural Slicing Using Dependence Graphs, *ACM Transactions on Programming Languages and Systems*, **12**(1), 1990, 26–60.

[9] Komarath, B., Sarma, J., Sunil, K. S.: On the Complexity of L-reachability, *Descriptional Complexity of Formal Systems - 16th International Workshop, DCFS 2014, Turku, Finland, August 5-8, 2014. Proceedings*, 2014.

[10] Reingold, O.: Undirected connectivity in log-space, *Journal of the ACM*, **55**(4), 2008.

[11] Reps, T. W.: On the Sequential Nature of Interprocedural Program-Analysis Problems, *Acta Informatica*, **33**(8), 1996, 739–757.

[12] Reps, T. W.: Program analysis via graph reachability, *Information & Software Technology*, **40**(11-12), 1998, 701–726.

[13] Sudborough, I. H.: A Note on Tape-Bounded Complexity Classes and Linear Context-Free languages, *Journal of the ACM*, **22**(4), 1975, 499–500.

[14] Sudborough, I. H.: On the Tape Complexity of Deterministic Context-Free Languages, *Journal of the ACM*, **25**(3), 1978, 405–414.

[15] Ullman, J. D., van Gelder, A.: Parallel Complexity of Logical Query Programs, *Algorithmica*, **3**, 1988, 5–42.

[16] Yannakakis, M.: Graph-Theoretic Methods in Database Theory, *Proceedings of the 9th ACM Symposium on Principles of Database Systems*, 1990.